

REMARKS/ARGUMENTS:

Applicant respectfully requests reconsideration of the present application in view of the below remarks.

Pending claims 28-56 are rejected, of which claims 28, 45, and 56 are in independent form.

The Prior Art Rejections

Claims 28-56 are rejected under 35 U.S.C. §103(e) over U.S. Patent No. 6,567,811 to Edwards (“Edwards”) in view of U.S. Patent Publication No. 2002/0078296 to Nakamura et al. (“Nakamura”).

Claim 28 requires a method for managing *data that may be replicated* from one or more *volumes of data* that are part of a *first volume group*, including:

discovering logical information related to the *one or more volumes of data that are part of the first volume group* on the first computer system and creating a map of the logical information to physical devices on the first computer system;

mounting a *duplicate* of the one or more *volumes of data* on a *second computer system* having a second operating system and using the map to create *a second volume group* that is *substantially identical* to the *first volume group*.

Applicants maintain that independent claim 28 (and all claims dependent therefrom, namely claims 29-44) is patentable over the cited art, taken either alone or in combination. As the Office is well aware, to establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim

limitations. Applicants maintain that the Office has not established a *prima facie* case of obviousness in the instant case.

First, Applicants contend that the cited references, taken alone or in combination, do not teach or suggest each and every limitation of claim 28. In claim 28 above, the “*data that may be replicated from one or more volumes of data*,” refers at least to data that is a mirror image or *replication* of one or more volumes of data. As explained in Applicants’ terminology descriptions on page 13, lines 5-15, a mirror image (also referred to as a mirrored disk) is a *copy* or *replication* of a source or standard volume. In one embodiment of the invention, as explained on page 13, lines 5-15, this copy or replication can be disk-based. “*Replication*” is a term well known to those of skill in the art. For example, “*replication*” is defined in the Microsoft Computer Dictionary (4<sup>th</sup> ed., 1999 at page 383, attached hereto as Exhibit A) as follows:

“**replication** n. In a distributed database management system, the process of copying the database (or parts of it) to the other parts of the network. Replication allows distributed database management systems to remain synchronized.

Thus, as those of skill in the art understand, a mirror image/replication is capable of operating differently than a mere static identical copy of a source, because the mirror image/replication is capable of continually “reflecting” back what is on the source, to stay synchronized with the source. For example, as described on page 13, lines 5-21 of Applicants’ specification, if data on the standard volume changes, the same changes are immediately applied to the mirrored disk, and mirrors can be synchronized in either direction (from standard disk to mirrored disk and vice versa). As further explained in Applicants’ specification at page 9, lines 16-19, such *replication* of a volume (mirroring) is useful because it [the mirror] can allow access to production volumes (and the data in them) even while backup operations are being performed on the mirror’s source, and the replication/mirrored volumes that were accessed can later be used to synchronize the mirror’s source. Thus, applying the invention of claim 28, a primary volume group and logical volumes and disks can all remain mounted and available to a source host with no outage of usage, while a replicated volume group with separate logical volumes and disks is built and could be automatically mounted on the same host or a remote host. This application would be impossible with Edwards, as is explained below.

For the “data that may be replicated” of claim 28 (i.e., the replication) and the rest of claim 28 to make sense, however, there has to be something to be replicated – and this is where Edwards differs significantly from the present invention. Edwards does not copy or replicate a volume group, while leaving the original intact. Rather, Edwards relates to merging volume groups, such as by moving the original data or filesystem to a larger volume group, e.g., by reorganizing or rearranging volume groups, whereby the original volume group no longer exists and the result is only the new, merged/reorganized volume group (see FIG. 8 and col. 10, line 40 through col. 11, line 6 of Edwards). In numerous examples and Figures, Edwards makes it a point to show that it is merging two volume groups, with one of the groups being removed. For example, Edwards describes situations such as moving data to a larger volume group (Edwards at col. 5, lines 34-35), rearrangement and combination of UNIX data structures (Edwards at col. 5, lines 48-49), system reconfigurations (Edwards at col. 10, lines 29-30) and reorganization (see comment text in FIG. 9A of Edwards). In Edwards, a filesystem is brought from a first “disappearing” volume group to a second “absorbing” volume group, where the second volume group actually absorbs the contents of the first volume group and where the first volume group is zeroed out or written over (see FIG. 8 and col. 10, lines 27-54 of Edwards).

As explained above, because Edwards merges a first volume group into a second volume group while removing the first volume group, the second volume group of Edwards is not a replication or mirror of the first volume group, but instead is used to replace it. As a consequence, claim 28’s requirement of “*one or more volumes of data that are [not “*were*”] part of a first volume group*”, from which “data may be replicated,” cannot possibly exist under Edwards, as the second volume of Edwards is not a replication of anything - it is, in fact, the new “first” source or original volume group. Thus, Edwards does not mount a duplicate, as required by claim 28. If the teachings of Edwards are followed, the original (or first) volume group, which per claim 28 needs to exist, no longer exists. If the teachings of Edwards are followed, claim 28’s recitation of “*mounting a duplicate of the one or more volumes of data on a second computer system having a second operating system and using the map to create a second volume group that is substantially identical to the first volume group*” never occurs.

Nakamura does not overcome any of the deficiencies of Edwards and likewise does not teach or suggest each and every element of claim 28, taken alone or in combination with Edwards. Nakamura relates to controlling a *paired volume recreation system*, where data is copied from a main (first) storage system to a remote (second) storage system, where the main and remote systems require consistency of data between them, and where the consistency must be maintained even if transmission of data is temporarily stopped or suspended (see Nakamura at paragraphs 8-11). For example, Nakamura seems to discuss using a cache based system to store pending data writes while data replication between systems is taking place. This is quite different than the invention of claim 28.

Nakamura further relates to reducing the time it takes to copy a group paired logical volumes, where some of the paired volumes are copied immediately and some are instead put in a “suspend” status and are not copied until a so-called “disable reason” (e.g., a data transfer unit that needs restoration) is removed from the logical volumes (see Nakamura at paragraphs 20 and 22). Nakamura mentions making a replication of the target volume before putting the volume into “suspend status” (Nakamura paragraph 22), but this replication seem to be provided only to restore data if a fault occurs during Nakamura’s “paired volume recreation”. Nakamura is silent as to how the replication is made, whether a map is created or used, and whether or how the replication is managed, merely noting in paragraph 27 that it uses known techniques for copying data either between disk subsystems or within the same disk subsystem.

Further, Nakamura does not teach or suggest claim 28’s recitation of creating a map of logical information related to one or more volumes of data that are part of a first volume group on a first computer system, mounting a duplicate of the one or more volumes of data on a second computer system, and using the map to create a second volume group that is substantially identical to the first volume group.

Applicants further contend that the neither the references themselves nor the knowledge generally available in the art would suggest or motivate the combination of Edwards with

Nakamura. As explained above, Edwards relates to merging volume groups, whereas Nakamura relates to a paired volume recreation system. As noted above, Nakamura relates to reducing the time it takes to copy a group paired logical volumes, where some of the paired volumes are copied immediately and some are instead put in a “suspend” status and are not copied until a so-called “disable reason” is removed from the logical volumes. As further noted above, Nakamura mentions making a replication of the target volume before putting the volume into “suspend status”, but is silent as to how the replication is made and/or managed.

Edwards has no need for Nakamura’s teachings regarding maintaining data consistency between first and second volumes or reducing the time to copy a paired volume, since under Edwards the first volume doesn’t even exist after the reorganization/merging: Edwards merges the first volume into a second volume and then zeroes out or writes over the first volume (see FIG. 8 and col. 10, lines 27-54 of Edwards, and claims 1-15 of Edwards). Thus, Edwards does not teach or suggest the utility or desirability of any modifications to its system that relate back to managing or controlling the first volume of data to keep it consistent with a second volume of data. Thus, Applicants maintain that there is no motivation or suggestion in Edwards to modify it with Nakamura’s teachings relating to controlling a first volume of data and the second volume of data to keep them consistent with each other.

Applicants also fail to see how one of skill in the art would find any motivation in the art to modify Edwards with Nakamura. The Office has suggested that the motivation to combine is because both Edwards and Nakamura allegedly deal with “*logical volume duplication based on the networking system*” to improve a remote computer’s accessibility for implementing mounting and remote copying. Applicants respectfully disagree. Edwards does not deal with “logical volume duplication” but instead deals with reorganizing logical volumes to merge a first logical volume into a second logical volume. Edwards does not duplicate a volume – Edwards simply moves it elsewhere. Nakamura, as noted above, deals with maintaining data consistency between paired volumes of data and with delaying copying of data if a data volume has a fault. The two problems are not similar, and it would not be obvious to one of skill in the art to look to modify

Edwards in any manner that required maintaining a pair of volumes with consistent data between them, when Edwards expressly would combine the pair into a single volume.

Referring again to the criteria for *prima facie* obviousness, Applicants further note that one of skill in the art would have no reasonable expectation of success if Edwards were combined with Nakamura. As noted above, in Edwards, the first volume is absorbed into the second volume after the reorganization/merging, so there is no longer the pair of volumes used in Nakamura. Thus Applicants do not see how one of skill in the art could expect a modification such as Nakamura to be able to work with Edwards (or vice versa).

In view of the above, Applicants submit that claim 28 is patentably distinct over Edwards and Nakamura, whether taken alone or in combination. Claims 29 to 44 depend from and thus include the limitations of claim 28. Thus, Applicants submit that claim 29-44 are patentably distinct over the cited references for at least the reasons discussed above in connection with claim 28.

Claims 45-55 and 56 were rejected on the same grounds as the rejections of claim 28 to 44. Accordingly, Applicants submit that claims 45-55 are patentable over the cited references, taken alone or in combination, for at least the reasons discussed above in connection with claim 28. Applicants thus respectfully request that the rejection of claims 28 through 56 under 35 U.S.C. 103(a) be withdrawn.

In view of the above, Applicants submit that Claims 28-56 and the entire case are in condition for allowance. A notice of allowance is respectfully requested.

The Examiner is respectfully invited to telephone the undersigning attorney if there are any questions regarding this Amendment or this application.

Applicant does not acquiesce to any assertion made by the Examiner not specifically addressed herein.

The Assistant Commissioner is hereby authorized to charge payment of any additional fees associated with this communication or credit any overpayment to Deposit Account No. 500845, including but not limited to, any charges for extensions of time under 37 C.F.R. §1.136.

Respectfully submitted,

Dated: 13 - April - 2006

DALY, CROWLEY, MOFFORD & DURKEE, LLP

By: 

Marianne M. Downing

Reg. No. 42,870

354A Turnpike Street - Suite 301A

Canton, MA 02021-2714

Tel.: (781) 401-9988, x22

Fax: (781) 401-9966

*mmd@dc-m.com*

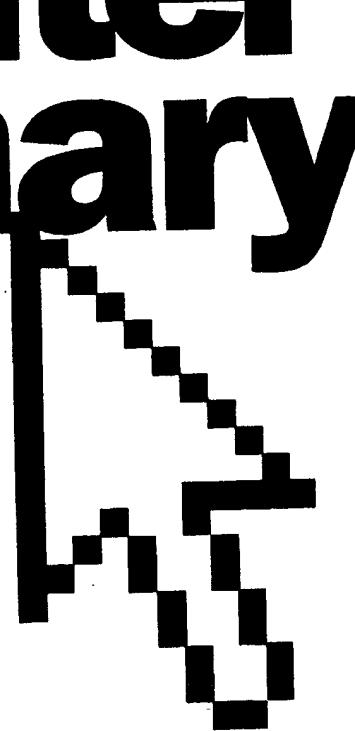
27328

**Microsoft®Press**

**EXHIBIT A**

Application No. 09/894,422  
Attny. Docket No. EMC-038PUS

**Microsoft®**  
**Computer**  
**Dictionary**  
**Fourth**  
**Edition**



**BEST AVAILABLE COPY**

## EXHIBIT A

Application No. 09/894,422  
Attny. Docket No. EMC-038PUS

### PUBLISHED BY

Microsoft Press  
A Division of Microsoft Corporation  
One Microsoft Way  
Redmond, Washington 98052-6399

Copyright © 1999 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

### Library of Congress Cataloging-in-Publication Data

Microsoft Computer Dictionary. -- 4th ed.

p. cm.

Previous eds. published under title: Microsoft Press computer dictionary

ISBN 0-7356-0615-3

1. Computers Dictionaries. 2. Microcomputers Dictionaries.

I. Microsoft Press computer dictionary.

QA76.15.M538 1999

004'.03--dc21

99-20168

CIP

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 MLML 4 3 2 1 0 9

Distributed in Canada by ITP Nelson, a division of Thomson Canada Limited.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at [mspress.microsoft.com](http://mspress.microsoft.com).

Macintosh, Power Macintosh, QuickTime, and TrueType fonts are registered trademarks of Apple Computer, Inc. Kodak is a registered trademark of the Eastman Kodak Company. Intel is a registered trademark and Indeo is a trademark of Intel Corporation. Active Desktop, Active Directory, ActiveMovie, Active Platform, ActiveX, Authenticode, BackOffice, DirectInput, DirectX, Microsoft, Microsoft Press, MS-DOS, MSN, NetMeeting, NetShow, Visual Basic, Visual C++, Visual J++, WebTV, WebTV Network, Win32, Win32s, Windows, Windows NT, and XENIX are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. PANTONE is a registered trademark of Pantone, Inc. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

Acquisitions Editor: Christey Bahn

Project Editor: Kim Fryer

**BEST AVAILABLE COPY**

## EXHIBIT A

Application No. 09/894,422  
Attny. Docket No. EMC-038PUS

repeating Ethernet *n.* See repeater.

repeat key *n.* On some keyboards, a key that must be held down at the same time as a character key to cause the character key's key code to be sent repeatedly. On most computer keyboards, however, a repeat key is not needed because a key automatically repeats if held down for longer than a brief delay. Compare typematic.

repeatKeys *n.* A feature of Windows 9x and Windows NT that allows a user to adjust or disable the typematic keyboard feature so as to accommodate users with restricted mobility, who may activate typematic by accident because they have trouble lifting their fingers from the keys. See also typematic. Compare BounceKeys, FilterKeys, MouseKeys, ShowSounds, SoundSentry, StickyKeys, ToggleKeys.

repetitive strain injury *n.* An occupational disorder of the tendons, ligaments, and nerves caused by the cumulative effects of prolonged repetitious movements. Repetitive strain injuries are appearing with increasing frequency among office workers who spend long hours typing at computerized workstations that are not equipped with safeguards such as wrist supports. Acronym: RSI. See also carpal tunnel syndrome, ergonomic keyboard, wrist support.

replace *vb.* To put new data in the place of other data, usually after conducting a search for the data to be replaced. Text-based applications such as word processors typically include search-and-replace commands. In such operations, both old and new data must be specified, and search-and-replace procedures may or may not be sensitive to uppercase and lowercase, depending on the application program. See also search<sup>1</sup>, search and replace.

replication *n.* In a distributed database management system, the process of copying the database (or parts of it) to the other parts of the network. Replication allows distributed database systems to remain synchronized. See also distributed database, distributed database management system.

report *n.* The presentation of information about a given topic, typically in printed form. Reports prepared with computers and appropriate software can include text, graphics, and charts. Database programs can include special software for creating report forms and generating reports. Desktop publishing software and laser printers or typesetting equipment can be used to produce publication-quality output.

report generator *n.* An application, commonly part of a database management program, that uses a report "form" created by the user to lay out and print the contents of a database. A report generator is used to select specific record fields or ranges of records, and to make the output attractive by including such features as headings, running heads, page numbers, and fonts.

report writer *n.* See report generator.

repository *n.* 1. A collection of information about a computing system. 2. A superset of a data dictionary. See also data dictionary.

reprogrammable PROM *n.* See EPROM.

reprogrammable read-only memory *n.* See EPROM.

Request for Comments *n.* See RFC.

Request for Discussion *n.* A formal proposal for a discussion concerning the addition of a newsgroup to the Usenet hierarchy, the first step in a process that ends with a call for votes. Acronym: RFD. See also traditional newsgroup hierarchy, Usenet.

Request to Send *n.* See RTS.

required hyphen *n.* See hyphen.

Research Libraries Information Network *n.* The combined online catalog of the Research Libraries Group, which includes many of the major research libraries in the United States. Acronym: RLIN.

reserve *n.* A command that allocates contiguous disk space for the device instance's workspace. Digital video devices recognize this command.

reserve accumulator *n.* An auxiliary storage register generally used to store the intermediate results of an extended calculation.

reserved character *n.* A keyboard character that has a special meaning to a program and, as a result, normally cannot be used in assigning names to files, documents, and other user-generated tools such as macros. Characters commonly reserved for special uses include the asterisk (\*), forward slash (/), backslash (\), question mark (?), and vertical bar (|).

reserved memory *n.* See UMA.

reserved word *n.* A word that has special meaning to a program or in a programming language. Reserved words usually include those used for control statements (IF, FOR, END), data declarations, and the like. A reserved word can be used only in certain pre-defined circumstances; it cannot be used in naming documents, files, labels, variables, or user-generated tools such as macros.

R